

IN THE SPECIFICATION

Please replace the paragraph beginning on page 92, line 1, with the following:

With reference to the flowchart of FIG. ~~[[16D]]~~ 16C, the steps involved in simulation job execution step **1627** of FIG. ~~[[16C]]~~ 16B are depicted in greater detail. The process of executing a simulation job on a simulation client begins with step **1631**, which depicts the simulation client obtaining a copy of the model corresponding to the given simulation job provided by the model servers. As illustrated at step **1638**, the simulation client communicates with instrumentation server **1699** to obtain and process control information for the instrumentation events within the simulation model. Proceeding to step **1632**, the simulation model is loaded into a hardware simulator or memory **44** of the simulation client.

Please replace the paragraph beginning on page 127 line 17 with the following.

It should be noted that a preliminarily non-redundant harvest event, as recorded within local harvest hit table **2201**, may in fact be redundant with respect to interim testcase activity within batch simulation farm **1601**. In the time interval between the initialization of local harvest hit table **2201** and the completion of the object testcase, another simulation client may detect and record the occurrence of one or more of the preliminarily non-redundant harvest events detected by simulator **1735**. In this case, harvest events appearing to be new in accordance with local harvest hit table **2201**, have actually already occurred during another testcase. In accordance with a preferred embodiment of the present invention, API entry point ~~rpt_hrv()~~ rpt_harv() **2202** may undertake additional processing steps to ensure that an apparently newly occurring harvest event has not actually occurred elsewhere within batch simulation farm **1601**.

Please replace the paragraph beginning on page 127 line 28 with the following.

In the direct mode of non-redundancy status inquiry, API entry point **2202** opens a direct network connection (e.g. Unix socket connection) on network **1720** to a harvest manager program **2215**, which executes on instrumentation server **1699**. Over this direct network connection, API entry point ~~rpt_hrv()~~ rpt_harv() **2202** delivers an aggregate instrumentation data

packet consisting of the contents of harvest cycle counters **422a-422n**, harvest flags **423a-423n**, and the name of the current testcase.

Please replace the paragraph beginning on page 129 line 28 with the following.

To this end, an indirect non-redundancy verification inquiry is utilized wherein API entry point **2202** uses only local harvest hit table **2201** in determining whether or not to collect a testcase. In indirect non-redundancy verification, API entry point ~~*rpt_hrv()*~~ *rpt_harv()* **2202** bypasses the step of further validating an apparently new harvest event with harvest manager **2215**. Communication overhead associated with a direct network connection with harvest manager program **2215** is thus reduced at the cost of potential redundancy in harvest testcase bucket **2300**.

Please replace the paragraph beginning on page 131 line 25 with the following.

However, in indirect non-redundancy verification mode, it is possible for multiple testcases to be delivered from simulation clients to harvest testcase bucket **2300** for the same harvest event. As one example, two simulation clients may receive the same harvest hit table content from the *init_harv()* call by API entry point **2200**. These simulation clients then independently execute differing testcases that both trigger the same, preliminarily non-redundant (in accordance with their respective local harvest hit tables) harvest event. In both simulation clients, API entry point ~~*rpt_hrv()*~~ *rpt_harv()* **2202** will instruct RTX **1702** to harvest their respective testcases.

Please replace the paragraph beginning on page 132 line 27 with the following.

As an additional optional ~~additional~~ processing step available in either direct or indirect harvest mode, API entry point **2202** may also communicate with instrumentation server **1699** and/or shared file system **1609** to obtain an updated copy of the harvest hit table. Typically, this data is obtained from shared file system **1609** to reduce the communication load on instrumentation server **1699**.

Please replace the paragraph beginning on page 134 line 1 with the following.

It should be noted that certain errors can occur that prevent RTX 1702 from successfully storing the current testcase on harvest testcase server 2210. However, when such an error occurs in direct non-redundancy verification mode, master harvest hit table 2205 has already been updated to indicate that the current testcase has been collected. Similarly, in indirect mode, master harvest hit table 2205 has been or will be similarly updated (barring errors in processing the aggregate instrumentation data packet sent to instrumentation server 1699 during step 2256 of FIG. 22B). Such failures to store the testcase can cause an inconsistency between master harvest hit table 2205 and harvest testcase bucket 2300 stored on ~~instrumentation~~ harvest testcase server 2210.

Please replace the paragraph beginning on page 135 line 4 with the following.

Proceeding to step 2254, wherein is depicted a determination of whether non-redundancy verification is to be performed in direct or indirect mode. If indirect mode is selected, and as illustrated at step 2256, the aggregate instrumentation packet generated in step 2253 is scheduled for a later delivery to harvest manager program 2215. If direct non-redundancy verification processing is selected by the simulation client, API routine ~~rpt_hrv()~~ rpt_harv() 2202 continues as shown at step 2255, by delivering the aggregate instrumentation data packet to harvest manager program 2215 via a direct network connection over network 1720 to validate the first occurrence status of the apparently new harvest events (step 2257).

Please replace the paragraph beginning on page 138 line 11 with the following.

The second source of inconsistency between the entries within master harvest hit table 2205 and the testcases recorded in harvest testcase bucket 2300, referred to hereafter as an “extraneous harvest testcase”, may occur as a result of the nature of the previously described indirect redundancy verification mode optionally undertaken by API entry point ~~rpt_hrv()~~ rpt_harv() 2202. As previously explained with reference to **Figures 21A-21C** 22A-22C, an indirect non-redundancy status inquiry results in the possibility that testcases which are redundant with respect to a given harvest event may be collected in association with a given simulation model within harvest testcase bucket 2300. This is despite the fact that only one testcase is recorded per harvest event in master harvest hit table 2205. The collection of potentially thousands of extraneous testcases within harvest testcase bucket 2300 is inherently

undesirable. Furthermore, it is not possible to definitively determine from master harvest hit table **2205** which harvest events are triggered by these extraneous testcases.

Please replace the paragraph beginning on page 140, line 22, with the following:

In non-hierarchical mode, harvest ~~events~~ testcases are collected in harvest testcase bucket **2300** without regard to each specific instance of a given harvest event. Once a testcase exercising a particular instance of a harvest event is detected and recorded, no further testcases are collected for that event. To this end, instance identifier field **2361** is ignored by harvest manager program **2215** when operating in non-hierarchical harvest mode.